

Sizing the BFGS and DFP Updates:
A Numerical Study

M. Contreras
and
R.A. Tapia

July, 1991

TR91-19

Report Documentation Page			Form Approved OMB No. 0704-0188		
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE JUL 1991		2. REPORT TYPE		3. DATES COVERED 00-00-1991 to 00-00-1991	
4. TITLE AND SUBTITLE Sizing the BFGS and DFP Updates: A Numerical Study			5a. CONTRACT NUMBER		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S)			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Computational and Applied Mathematics Department ,Rice University,6100 Main Street MS 134,Houston,TX,77005-1892			8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSOR/MONITOR'S ACRONYM(S)		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S)		
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 15	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

Sizing the BFGS and DFP Updates: A Numerical Study*

Martha Contreras [†] Richard A. Tapia [‡]

Abstract

In this study we develop and test a strategy for selectively sizing (multiplying by an appropriate scalar) the approximate Hessian matrix before it is updated in the BFGS and DFP trust-region methods for unconstrained optimization. Our numerical results imply that for use with the DFP update the Oren-Luenberger sizing factor is completely satisfactory and selective sizing is vastly superior to the alternatives of never sizing, or first-iteration sizing, and is slightly better than the alternative of always sizing. Numerical experimentation showed that the Oren-Luenberger sizing factor is not a satisfactory sizing factor for use with the BFGS update. Therefore, based on our newly acquired understanding of the situation, we propose a damped Oren-Luenberger sizing factor to be used with the BFGS update. Our numerical experimentation implies that selectively sizing the BFGS update with the damped Oren-Luenberger sizing factor is superior to the alternatives. These results contradict the folk-axiom that sizing should be done only at the first iteration. They also show that without sufficient sizing, DFP is vastly inferior to BFGS; however, when selectively sized, DFP is competitive with BFGS.

*Research supported in part by NSF Coop. Agr. No. CCR-8809615, AFOSR 89-0363 and DOE DEFG05-86ER25017, and AR0 9DAAL03-90-G-0093. This paper was presented at the ICIAM 91 International Conference, Washington D.C., July 1991

[†]This author is currently a graduate student in the Mathematics Department, University of California Riverside, Riverside, California 92521, and has been a part of the Summer Student Visitor Program at the Center for Research and Parallel Computation at Rice University for the past two years.

[‡]Department of Mathematical Sciences and Center for Research in Parallel Computation, Rice University, Houston, Texas 77251-1892.

1 Introduction

In this note we study the effect that sizing (multiplying the approximate Hessian by an appropriate scalar before it is updated) has on the performance of the BFGS and DFP trust-region methods for unconstrained optimization. We suggest sizing strategies for both these updates and include considerable numerical experimentation that indicates that our selective sizing is superior to the alternatives of never sizing, sizing only at the first iteration, or sizing at every iteration. Our experimentation also supports the common belief that sizing is more critical when using the DFP update than it is when using the BFGS update. In fact we were pleasantly surprised to see that in our experiments, when the DFP algorithm was sized according to our strategy, it performed numerically as well as the BFGS algorithm.

In the remainder of this section we present some history and preliminaries concerning the notion of sizing. In Section 2 we consider some interesting properties of the Oren-Luenberger sizing factor. Our numerical studies showed that the Oren-Luenberger sizing factor gives satisfactory performance when used with the DFP update, but not with the BFGS update. Hence, in Section 3, guided by the understanding gained in Section 2 and our numerical experimentation, we motivate and propose the damped Oren-Luenberger sizing factor. In Section 4 we describe our selective sizing strategy. Our numerical results are presented in Section 5, and in Section 6 we make some concluding remarks. We make the basic assumption that the reader is familiar with at least chapters 6 and 9 of Dennis and Schnabel [6].

In 1974 Oren and Luenberger [12] proposed a class of secant methods which they referred to as self-scaling variable metric (SSVM) methods. Soon after, Oren and Spedicato [13] identified a subclass of the SSVM methods that had various desirable properties. Shanno and Phua [15], among other things, studied the BFGS secant update and its associated scaling as a member of the Oren-Spedicato subclass of SSVM methods. They argued that the BFGS update should be scaled only at the first iteration, as opposed to every iteration as is implied by the self-scaling philosophy. They presented numerical evidence that showed that in general scaling only the initial iteration was superior to scaling at every iteration. Our numerical experiments corroborated their findings when the Oren-Luenberger sizing factor was used; however, the opposite was true when the centered Oren-Luenberger sizing factor was used. In 1981 Dennis, Gay and Welsch [4] proposed the now well-known NL2SOL algorithm for the nonlinear least-squares problem. They incorporated several novel features into their algorithm. The basic framework consists of a Gauss-Newton trust-region method. In order to handle large residual problems they maintain a structured BFGS secant approximation to S_k , the second-order part of the least-squares Hessian, and adaptively decide when to use this approximation. One of their major concerns was that since secant methods do not generate approximations that become arbitrarily accurate as the iteration proceeds, the approximation to this second-order information may impede the fast convergence for zero residual problems, i.e., the second-order information may be zero at the solution, but the approximations S_k may not converge to zero. This deficiency was overcome in the following clever manner. First, they observed that the Oren-Luenberger scaling for the BFGS update, say γ_k , had the property that the interval spectrum of $\gamma_k S_k$ essentially overlapped the interval spectrum of the matrix of the exact second-order information. Hence, by scaling by γ_k one could guarantee that $\gamma_k S_k$ would converge to zero when the second-order information was

zero at the solution. In this application it was critical that the scaling was done at each iteration; unlike the Shanno-Phua applications. Very recently, Huschens [8] utilizing the Dennis-Martinez-Tapia BFGS structure principle [5], has given a most elegant solution to this problem. In order to emphasize this spectrum shifting property, Dennis, Gay and Welsch decided to call the process of multiplying S_k by γ_k “sizing” as opposed to scaling. We will follow Dennis, Gay and Welsch in this usage and actually rigorously define what we perceive this notion to be. We begin with the following definitions.

Definition 1.1 *By the convex spectrum of a collection of $n \times n$ matrices A_1, \dots, A_m denoted $\text{conspectrum}(A_1, \dots, A_m)$, we mean the convex hull of the eigenvalues of A_1, \dots, A_m .*

When A is symmetric the convex spectrum of A is an interval, and we therefore also refer to it as the interval spectrum of A .

Definition 1.2 *We say that the scalar γ sizes $B \in R^{n \times n}$ relative to $A_1, \dots, A_m \in R^{n \times n}$ if $\text{conspectrum}(\gamma B) \cap \text{conspectrum}(A_1, \dots, A_m) \neq \emptyset$.*

The following proposition illustrates the notion of sizing. It is easily extended to the general case:

Proposition 1.1 *Let $A, B \in R^{n \times n}$ be symmetric matrices. Then the scalar γ sizes B relative to A if and only if there exists $u, v \in R^n$ such that*

$$\gamma = \frac{u^T A u}{u^T u} \frac{v^T v}{v^T B v} \quad (1.1)$$

Proof. The proof follows directly from well-known properties of the Rayleigh quotient. \square

Our objective is to size the approximate Hessian relative to the true Hessian before the update is made in a secant method. Our formal definition of this notion requires some motivation.

Recall that by a secant method for approximating the minimizer of the C^2 function $f : R^n \rightarrow R$ we mean the iterative procedure

$$x_{k+1} = x_k - B_k^{-1} \nabla f(x_k), \quad k = 0, 1, \dots \quad (1.2)$$

where B_k is an approximation to $\nabla^2 f(x_k)$ and $B_{k+1} = U(s, y, B_k)$, the update of B_k , satisfies the secant equation

$$B_{k+1} s = y \quad (1.3)$$

with $s = x_{k+1} - x_k$ and $y = \nabla f(x_{k+1}) - \nabla f(x_k)$.

Without evaluating additional functions or gradients, the only Hessian information we have is y , and this is only approximate information. Specifically, y is an approximation to $\nabla^2 f(x)s$. Moreover, since y is a vector and the standard mean-value-theorem does not necessarily hold for vector-valued functions, we do not know that there exists $\hat{x} \in R^n$ such that $y = \nabla^2 f(\hat{x})s$. However, McLeod’s mean-value theorem [9] for vector-valued functions says that there exist $x_1, x_2, \dots, x_n \in R^n$ such that y is contained in the convex hull of $\nabla^2 f(x_1)s, \nabla^2 f(x_2)s, \dots, \nabla^2 f(x_n)s$. Hence, when our Hessian information comes only from y , it would be shortsighted of us to define sizing in terms of the Hessian of f at just one point. We therefore propose the following definition.

Definition 1.3 Consider a C^2 function $f : R^n \rightarrow R$. We say that the scalar γ sizes $B \in R^{n \times n}$ relative to the Hessian of f if there exists $x_1, x_2, \dots, x_m \in R^n$ such that

$$\text{conspectrum}(\gamma B) \cap \text{conspectrum}(\nabla^2 f(x_1), \nabla^2 f(x_2), \dots, \nabla^2 f(x_m)) \neq \phi. \quad (1.4)$$

We call the integer m the degree of the sizing factor γ .

The Oren-Luenberger sizing factor associated with the BFGS update and used by Shanno and Phua and Dennis, Gay, and Welsch is

$$\gamma = \frac{y^T s}{s^T B s}. \quad (1.5)$$

By the mean-value theorem we can write $\gamma = \frac{s^T \nabla^2 f(x + \theta s) s}{s^T B s}$ for some $0 < \theta < 1$; hence γ sizes B relative to the Hessian of f , and it is of degree 1.

It is worth noting that as a direct consequence of the secant equation (1.3) we have

$$\frac{y^T s}{s^T B_{k+1} s} = 1; \quad (1.6)$$

hence in any secant method all approximate Hessian approximations, except the initial approximation, are automatically sized relative to the Hessian of f . This phenomenon seemingly adds credibility to the Shanno-Phua doctrine of sizing only at the initial iteration.

Dennis and Wolkowicz [7], defer to Shanno and Phua [15], and consider some interesting alternatives to sizing after the first step.

Carter [3] in a very interesting work presented three procedures for safeguarding Hessian approximations. All three procedures selectively made changes to the Hessian approximations. One of these procedures (the one he liked least) is quite similar to ours. For the BFGS update he uses as the sizing factor for B_k

$$\gamma = \frac{y^T s}{s^T s} \frac{g_+^T g_+}{g_+^T B_k g_+} \quad (1.7)$$

where $s = x_{k+1} - x_k$, $y = \nabla f(x_{k+1}) - \nabla f(x_k)$ and $g_+ = \nabla f(x_{k+1})$. Moreover, he sizes whenever γ in (1.7) is less than 1. He stated that this choice of sizing for the BFGS method performed relatively poorly on large dimensional problems. We experienced a similar phenomenon with the Oren-Luenberger sizing factor (1.5); more will be said about this issue in Section 3. Observe that Carter's choice (1.7) sizes B_k relative to the Hessian of f and it is of degree 1.

Al-Baali [1] also considered selective scaling for the BFGS update using various scaling strategies.

2 Oren-Luenberger Sizing

The following proposition and its corollary played a major motivational role in our original plan to selectively size both the DFP and the BFGS update with the Oren-Luenberger sizing factor (1.5). In the proposition we assume the standard assumptions for secant method

theory, see Dennis and Schnabel [6] for details, and we consider a secant method of the form (1.2)–(1.3) which generates symmetric and positive definite updates, e.g. the DFP or the BFGS secant method.

Proposition 2.1 *If the secant method (1.2)–(1.3) with steplength 1 converges q -superlinearly, then the Oren-Luenberger sizing factory $\gamma_k = \frac{y_k^T}{s_k^T B_k s_k}$ converges to 1.*

Proof: Suppose that $\{x_k\}$ converges to the minimizer x_* . From the standard assumptions we know that $\nabla^2 f(x_*)$ is positive definite. Let $\lambda^* > 0$ be the smallest, eigenvalue of $\nabla^2 f(x_*)$. In what follows the quantities s, y , and B should all be viewed with a subscript k . We have

$$\left| \frac{y^T s}{s^T s} - \frac{s^T B s}{s^T s} \right| \leq \frac{\|y - B s\|}{\|s\|} = \frac{\|\nabla^2 f(x_*) s - B s\|}{\|s\|} + O(\|s\|). \quad (2.1)$$

Also,

$$\left\| \frac{s^T \nabla^2 f(x_*) s}{s^T s} - \frac{s^T B s}{s^T s} \right\| \leq \frac{\|\nabla^2 f(x_*) s - B s\|}{\|s\|}. \quad (2.2)$$

The Dennis Moré characterization of q -superlinear convergence (Theorem 8.2.4 of Dennis and Schnabel [6]) implies that the right-hand sides of (2.1) and (2.2) converge to zero. From (2.2) and properties of the Rayleigh quotient we have that $\liminf \frac{s^T B s}{s^T s} \geq \lambda^*$. Let us write

$$\frac{y^T s}{s^T B s} - 1 = \frac{s^T s}{s^T B s} \left(\frac{y^T s}{s^T s} - \frac{s^T B s}{s^T s} \right). \quad (2.3)$$

The proof now follows from (2.1) and (2.3) since we have established that $\frac{s^T s}{s^T B s}$ is uniformly bounded in k . \square

Corollary 2.1 *Proposition 2.1 holds for the BFGS secant method without the assumption that the convergence is q -superlinear.*

Proof: Very recently Byrd, Tapia and Zhang [2] have demonstrated that under the standard assumptions if the BFGS secant method with steplength 1 converges, then the convergence is q -superlinear. \square

Armed with this encouragement we initially set out to selectively size both the BFGS and DFP secant trust-region methods using the Oren-Luenberger sizing factor. From the very beginning our success for DFP was immediate and significant. Even in the case when we sized at every iteration, the sizing factor seemed to be converging to 1. Initially we also experienced good success for the BFGS trust-region method. Only when we tried large dimensional problems, e.g. problems of dimension greater than 20, did we conclude that the Oren-Luenberger sizing and the bFGS secant update were a bad fit. For these larger problems sizing often hurt the performance; yet it often helped the performance, we could not determine when it would help or when it would hurt. It is interesting to note that Carter [3] mentioned that he also experienced poor performance for the larger dimensional problems when he selectively applied his sizing factor (1.7) to the BFGS secant method. We noticed that in several examples the Oren-Luenberger sizing factor did not converge to 1. However,

in these examples the trust-region was active in some of the final iterations; hence Corollary 2.1 did not seem to be violated.

In a very interesting recent work Nocedal and Yuan [11] have studied the asymptotic behavior of the BFGS secant method using Oren-Luenberger sizing. Their theory implies that the sizing factor need not converge to 1 even if steplength 1 is taken. Hence, according to Proposition 2.1, Oren-Luenberger sizing may impede superlinear convergence. Whether such a theory exists for the DFP secant update is perhaps an interesting open question.

3 Centered Oren-Luenberger Sizing

We have argued that Oren-Luenberger sizing (1.5) and the BFGS update are not an effective combination. We also feel that the same is probably true for Carter (1.7) sizing and the BFGS update. Since in trust-region methods one does not have access to B^{-1} we do not have the option of using so-called “inverse sizing”, i.e., multiplying B by $\gamma = \frac{y^T B^{-1} y}{y^T s}$. It is of some interest to note that according to our formal notion of sizing (Definition 1.3), “inverse sizing” is not a sizing of B .

At this juncture it is our considered opinion that for the BFGS trust-region method there is no truly effective sizing factor of degree 1. Essentially we believe that a sizing factor of degree 1 does not carry enough information to consistently improve the overlap of the respective spectra for large dimensional problems. Therefore, in the following we will attempt to construct an effective sizing factor of degree greater than 1 for the BFGS update.

Let us recall that the objective of sizing is to overlap the respective spectra. An ideal overlap would match the center of the interval spectrum of B with the center of the interval spectrum of $\nabla^2 f(x)$. Of course the centers of these interval spectra are not known. However, given various points in the interval spectrum, a convex combination, in particular the arithmetic mean (average), of these points serves as an approximation to the center of the interval spectrum. This is the central idea behind the sizing factor we are about to construct.

In what follows we use the notation $s = x_{k+1} - x_k$, $s_- = x_k - x_{k-1}$, $y = \nabla f(x_{k+1}) - \nabla f(x_k)$, $y_- = \nabla f(x_k) - \nabla f(x_{k-1})$. By the *centered Oren-Luenberger sizing factor* we mean

$$\hat{\gamma}(\theta_k) = \frac{(1 - \theta_k)(y_-^T s_-)/(s_-^T s_-) + (\theta_k)(y^T s)/(s^T s)}{(1 - \theta_k)(s_-^T B_k s_-)/(s_-^T s_-) + (\theta_k)(s^T B_k s)/(s^T s)} \quad (3.1)$$

where $0 \leq \theta_k \leq 1$. Clearly for $\theta_k = 1$ we obtain the Oren-Luenberger sizing factor, and for $\theta_k = 0$ we obtain the constant 1. To see this recall the secant equation $B_k s_- = y_-$. Our intuition tells us that we should be particularly interested in $\hat{\gamma}(\theta_k)$ for $\theta_k = 1/2$.

The following proposition shows that $\hat{\gamma}(\theta_k)$ is also a “softening” or “dampening” of the Oren-Luenberger factor γ_k . Recall that in standard implementations of the DFP and BFGS secant updates we don’t update if $y^T s \leq 0$, and we are guaranteed that $y^T s > 0$ near the solution. These concerns guarantee that B_k will be positive definite. Hence, no generality is lost by assuming that $y^T s$, $y_-^T s_-$, and $s^T B_k s$ are all positive for the purpose of the following proposition.

Proposition 3.1 *The following statements are equivalent:*

- (i) $\gamma_k < 1$;

- (ii) $\hat{\gamma}(\theta_k) < 1$;
- (iii) $\gamma_k < \hat{\gamma}(\theta_k)$.

Proof: The proof is straightforward once we recall that $B_k s_- = y_-$. \square

Proposition 3.2 *The scalar $\hat{\gamma}(\theta_k)$ sizes B_k relative to the Hessian of f and is of degree 2.*

Proof: The proof is straightforward. \square

It should be clear that a centered Oren-Luenberger sizing factor of degree m could be defined analogously to (3.1) by employing the quantities $s_k, s_{k-1}, \dots, s_{k-m+1}$ and $y_k, y_{k-1}, \dots, y_{k-m+1}$. Since the centered Oren-Luenberger sizing factor of degree 2 gave satisfactory numerical results we did not experiment with those of larger degree.

4 Selective Sizing Strategy

Our strategies, for both the BFGS and DFP, are to size at the first iteration using the Oren-Luenberger sizing factor (1.5), and then selectively size at other iterations using the Oren-Luenberger for the BFGS update. Recall that $\hat{\gamma}(\theta_k)$ denotes the centered Oren-Luenberger sizing factor (3.1) and γ_k denotes the Oren-Luenberger sizing factor $\frac{y^T s}{s^T B s}$. Also ϵ_1 and ϵ_2 are small positive constants and τ_1 and τ_2 are nonnegative constants.

General Sizing Strategy

For $k = 0$ size B_0 by $\max(\epsilon_2, \gamma_0)$.

For $k = 1, 2, \dots$ let $\theta_k = \min(\tau_1, \tau_2 \|s_k\|)$.

If $\hat{\gamma}(\theta_k) \leq 1 - \epsilon_1$, then size B_k by $\max(\epsilon_2, \hat{\gamma}(\theta_k))$.

DFP Sizing Strategy

Choose $\tau_1 = 1$ and τ_2 large.

BFGS Sizing Strategy

Choose $\tau_1 = 1/2$ and τ_2 large.

The inclusion of the features ϵ_1, ϵ_2 and τ_2 serve as safeguards for our sizing strategy. The use of ϵ_2 prevents us from inadvertently creating a nearly singular matrix by sizing with an excessively small constant. On rare occasions this feature was selected in our numerical experiments. The use of $\epsilon_1 > 0$ and $\tau_2 > 0$ ensures that sizing will eventually be shut off. To see this observe that with $\tau_2 > 0$ we have that $\theta_k \rightarrow 0$, so $\hat{\gamma}(\theta_k) \rightarrow 1$, i.e., we have forced the sizing factor to converge to 1; in contrast to depending on, as yet, not fully understood theory. However, in our numerical experiments described in the following section we did not use this feature and always sized with $\theta_k = 1$ for DFP and $\theta_k = 1/2$ for BFGS.

The choice $\tau_1 = 1$ is extremely important for the DFP method. It has been our experience that DFP loves to be sized by the Oren-Luenberger sizing factor. On the other hand the choice $\tau_1 = 1/2$ seems to be extremely important for the BFGS method. It seems as if sizing the BFGS method is a delicate issue especially for large dimensional problems, and using 2 pieces of approximate Hessian information is significantly better than using just 1.

5 Numerical Results

The sizing procedures of Section 4 were adapted to the secant trust-region algorithm in the Dennis-Schnabel code [6]. This code uses the locally constrained optimal step, or “Hook step”, of Hebden (1973) and Moré [10] to obtain an approximate solution to the trust-region subproblem. The code allows only the BFGS secant update; so we added the option of using the DFP secant update.

Six test functions plus Oren’s Power Function were selected from the standard test set of Moré, Garbow and Hillstom [10]. In the following table we use the notation $f(d)$. Here f is an integer running from 1 to 7 and d denotes the number of variables or function arguments. We use $f = 1$ to signify the Helical Valley Function, $f = 2$ signifies the Penalty Function, $f = 3$ signifies the Extended Powell Singular Function, $f = 4$ signifies Oren’s Power Function, $f = 5$ signifies the Extended Rosenbrock Function, $f = 6$ signifies the Trigonometric Function, and finally $f = 7$ signifies the Wood Function. All of the above functions can be found in Moré, Garbow, and Hillstom [10], except for Oren’s Power Function, and this latter function is

$$f(x) = (x^T A x)^2$$

where $A = [a_{ij}]$, with $a_{ij} = i\delta_{ij}$ and δ_{ij} denotes the Dirac delta function.

To test the effectiveness of our sizing strategy, we modified the Dennis-Schnabel optimization code to accept an arbitrary initial matrix B_0 , in its unmodified form it always takes $B_0 = I$ and sizes only at the first iteration. We chose several very ill-conditioned initial matrices as well as $B_0 = I$. The notation $D(a)$ denotes the diagonal matrix with all diagonal elements set equal to a . The notation $D(a, b)$ means the diagonal matrix which alternates a and b along its diagonal. Initial matrix BTZ, $q = n$ means the Byrd-Tapia-Zhang matrix, $B_0 = \text{diagonal}(\alpha_1, \dots, \alpha_n) + I$, where $\alpha_i = (1 - (\frac{i-1}{n-1}))(10^q - 1)$. The starting point is denoted by x_0 . We have chosen several nonstandard starting points, as well as the standard starting points that can also be found in Moré, Garbow, and Hillstom. In all our experiments we allowed a maximum of 300 iterations. The letter F means failure, i.e., lack of convergence in 300 iterations. The stopping criterion used is the one used by the Dennis-Schnabel optimization code, i.e., we define convergence when the relative gradient is less than or equal to 1.0×10^{-5} . The machine used to obtain these numerical results was a Sun 3/160.

In Table 2, the letters “n.s.” mean we never size B_k before updating; “a.s.” means we always size B_k when updating; “f.i.s.” means we size only on the first iteration. “S.S(n.l.s)” means two things: the S.S. part denotes the number of iterations it took the algorithm to converge when we sized B_k as described in Section 4, and the (n.l.s) part means out of the number of iterations it took to converge, how many times we sized. The symbol F^* means the algorithm stopped very close to the answer after 300 iterations. Finally $(x_1, x_2, x_3, x_4, \dots)^{**}$ means repeat this set of numbers until the dimension of the problem.

We found that when updating B_k using the BFGS update effective choices for ϵ_1 and ϵ_2 were $\epsilon_1 = .05$ and $\epsilon_2 = .1$, and for the DFP update effective choices were $\epsilon_1 = .001$ and $\epsilon_2 = .1$. Our results were not particularly sensitive to these choices. Our choice for τ_2 was 1.0×10^6 ; hence our shut-off feature was never activated and we selectively sized until convergence. Table 1 is self explanatory and summarizes our numerical results for the various approaches applied to our 27 problem configurations.

Update	Never Sizing	Always Sizing	1 st Iteration Sizing	Selective Sizing
BFGS	12/27	20/27	15/27	27/27
DFP	2/27	24/27	0/27	26/27

TABLE 1: Summary of Success Ratios

6 Summary and Concluding Remarks

In this note we studied the effect that sizing the approximate Hessian has on the BFGS and DFP trust-region methods. We accomplished this objective by adapting the Dennis-Schnabel BFGS trust-region code [6] to accept a user supplied initial Hessian approximation (they always use the identity) and to allow the user the option of the DFP secant update as well as the BFGS. We then ran the code on several standard test problems from the optimization literature using various starting points and various initial Hessian approximations. In order to seriously challenge the algorithms studied we considered many nonstandard starting points which were quite far from the solution (as well as the standard starting points) and several extremely ill-conditioned initial Hessian approximations. While our initial Hessian approximations were all diagonal matrices, some had condition numbers as large as 10^{24} . This numerical study led us to the following opinions.

The DFP update desperately needs to be sized and the Oren-Luenberger sizing factor (1.5) is an excellent sizing factor for this update. In our study sizing at each iteration was only slightly inferior to our selective sizing strategy, and at this point we are not prepared to recommend one over the other. Without sizing, the DFP update is vastly inferior to the BFGS update. However, when properly sized it is superior to the unsized BFGS update and competitive with the selectively sized BFGS update. Our numerical experience has led us to the intuition that when sizing is working, the sizing factor is converging to 1. Moreover, always sizing seemed to demonstrate this behavior somewhat more consistently than our selective sizing did for the DFP update.

We learned that sizing the BFGS update is a delicate issue for large dimensional problems. Our experience led us to the belief that sizing the BFGS update with a factor based on approximate Hessian information only at one point would not be completely effective. For this reason we developed what we call the centered Oren-Luenberger sizing factor to be used with the BFGS update. It uses approximate Hessian information gained at the current point and at the previous point to approximate the centers of the respective interval spectra, and both intuitively and experimentally seems to do a better job of overlapping the spectrum of the approximate Hessian with the spectrum of the Hessian. While more research and experimentation is needed in this direction, we believe that we have presented a convincing case that it is a viable alternative to traditional sizing. Indeed with this new form of sizing we obtained a success rate of 100% for the problems we tried.

For both the BFGS and the DFP trust-region methods we found that, in general, we should size when our sizing factor was less than 1, and not size when it was greater than 1. These findings are clearly reflected in our sizing strategy described in Section 4, and we offer

$f(d)$	x_o	B_o	Upd	n.s.	a.s.	f.i.s.	S.S.(n.t.s)
1(3)	(-10,1,10)	$D(10^{12}, 10^{-12})$	bfgs	F	F	F	78(11)
			dfp	F	54	F	64(36)
1(3)	(-1000,0,0)	BTZ, $q = 12$	bfgs	F	78	F	82(20)
			dfp	F	92	F	85(43)
2(4)	(1,2,3,4)	I	bfgs	170	F	167	166(62)
			dfp	F	F	F	235(142)
2(4)	(-1000,2000,3,4)	I	bfgs	F	F	F	192(53)
			dfp	F	275	F	222(94)
2(20)	(1,2,3,...,20)	BTZ, $q = 12$	bfgs	F	F	F	104(49)
			dfp	F	78	F	93(62)
2(20)	(1,2,3,...,20)	I	bfgs	117	177	138	145(55)
			dfp	F	265	F	188(114)
3(4)	(30000,-10,0,10)	I	bfgs	98	108	153	91(62)
			dfp	F	95	F	98(70)
3(4)	(30000,-10,0,10)	$D(10^{12}, 10^{-12})$	bfgs	F	95	F	84(50)
			dfp	F	101	F	93(68)
3(24)	(30000,-10,0,10,...)**	I	bfgs	F	231	F	187(77)
			dfp	F	261	F	197(73)
3(24)	(3,-1,0,1,...)**	I	bfgs	61	68	49	48(20)
			dfp	F	63	F	83(29)
3(24)	(30000,-10,0,10,...)**	$D(10^{12}, 10^{-12})$	bfgs	F	F	F	F^*
			dfp	F	F	F	241(102)
4(4)	(80000,1,1,1)	I	bfgs	100	74	F	75(64)
			dfp	F	53	F	53(52)
4(4)	(1,1,1,1)	I	bfgs	35	26	56	26(25)
			dfp	168	19	F	19(18)
4(20)	(1,1,1,...,1)	I	bfgs	145	40	232	40(39)
			dfp	F	26	F	26(25)
4(20)	(1,1,1,...,1)	$D(10^{12}, 10^{-12})$	bfgs	F	215	F	155(44)
			dfp	F	82	F	89(64)
5(4)	(-1.2,1,-1.2,1)	I	bfgs	51	64	147	55(18)
			dfp	F	85	F	63(32)
5(4)	(-40000,1,-1.2,1)	BTZ, $q = 12$	bfgs	F	187	F	132(61)
			dfp	F	137	F	123(77)
5(20)	(-40000,1,-1.2,1)**	BTZ, $q = 12$	bfgs	F	F	F	244(72)
			dfp	F	203	F	172(81)
5(20)	(-100.5,40,-100.5,40,...)**	I	bfgs	F	210	243	186(50)
			dfp	F	296	F	263(86)
5(20)	(-1.2,1,-1.2,1)**	$D(10^{12}, 10^{-12})$	bfgs	F	183	172	104(18)
			dfp	F	170	F	170(58)
5(20)	(-40,-20,20,20,...)**	$D(10^8, 10^{-8})$	bfgs	F	F	215	F^*
			dfp	F	F	F	F^*
5(20)	(-1.2,1,-1.2,1,...)**	I	bfgs	136	84	44	135(20)
			dfp	F	146	F	125(38)
6(4)	(.5,.5,.58,.5)	BTZ, $q = 11$	bfgs	F	F	47	68(11)
			dfp	F	70	F	59(24)
6(30)	(.5,.5,.58,5,...)**	BTZ, $q = 12$	bfgs	F	122	F	129(46)
			dfp	F	86	F	84(40)
6(33)	(.5,.5,.58,.5,...)**	I	bfgs	39	34	138	34(15)
			dfp	191	52	F	36(10)
7(4)	(-3,-1,-3,-1)	I	bfgs	50	133	38	115(28)
			dfp	F	178	F	121(68)
7(4)	(-300,-100,-300,-100)	$D(99999)$	bfgs	100	71	152	75(39)
			dfp	F	135	F	90(58)

TABLE 2: Numerical Results

the following plausible explanation.

When the sizing factor is greater than 1, B_k is “small” relative to the Hessian in some (as yet undefined) sense; hence $s_k = -B_k^{-1}\nabla f(x_k)$ will be “large”. However, the trust region strategy can compensate for a “large” s_k ; it merely reduces the trust-region radius and resolves the subproblem. On the other hand when our sizing factor is less than 1, B_k is “large” in this undefined sense; hence s_k is “small”. The trust-region strategy accepts s_k , since it is “small” and contained in the trust region. This could continue for several iterations; producing very slow progress towards convergence, since the steps would be excessively “small”. In this case sizing would help by making B_{k+1} “smaller” and the resulting s_{k+1} “larger”.

We believe that it is appropriate to end this section with a discussion which demonstrates the value, the beauty, and the subtlety of sizing. Let us consider the application of the gradient method to the strictly convex quadratic functions $f(x)$ with Hessian A . The gradient method with steplength α_k is the iterative procedure

$$x_{k+1} = x_k - \alpha_k \nabla f(x_k). \quad (6.1)$$

By rewriting (6.1) as

$$x_{k+1} = x_k - (1/\alpha_k I)^{-1} \nabla f(x_k), \quad (6.2)$$

we can view the gradient method with steplength α_k as a quasi-Newton method which employs steplength 1, but sizes the approximate Hessian (the identity matrix) with the sizing factor $\gamma_k = \alpha_k^{-1}$. The standard gradient method chooses the steplength by a 1-dimensional minimization and in this case would give $\gamma_k = g_k^T A g_k / g_k^T g_k$ where $g_k = \nabla f(x_k)$. This is clearly a sizing factor for I and has some of the flavor of the Carter sizing factor (1.7). Recently, in a very elegant and interesting work Raydan [14] demonstrated that if instead one chooses the steplength by choosing the sizing factor $\gamma_{k+1} = \frac{s_k^T A s_k}{s_k^T s_k}$, where $s_k = x_{k+1} - x_k$, then superior convergence behavior is obtained and superlinear convergence results for a subclass of problems. Notice that this choice of steplength is equivalent to applying Oren-Luenberger sizing. Raydan’s theory demonstrates that while two Rayleigh quotients may both give points in the interval spectrum of A , one may have hidden properties which make it particularly effective for the application in question.

Acknowledgements

The authors thank all those individuals that took part in the lively discussions concerning this material following the ICIAM 91 presentation. These discussions influenced the current version of the paper. The first author also thanks Cristina Maciel for assistance and support during the earlier stages of the research.

References

- [1] M. AL-BAALI. Partial self scaling variable metric algorithms. Report 80, Dipartimento di Sistemi, Universita della Calabria, 1988.

- [2] R.H. BYRD, R.A. TAPIA, and Y.ZHANG. An SQP augmented Lagrangian BFGS algorithm for constrained optimization. Report TR89-4, Department of Mathematical Sciences, Rice University, Houston, Tx 77251-1892, 1989, to appear in *SIAM J. Opt.*
- [3] R. CARTER. Safeguarding Hessian approximations in trust region algorithms. Report TR87-12, Department of Mathematical Sciences, Rice University, Houston, Tx 77251-1892, 1987.
- [4] J.E. DENNIS Jr., D.GAY, and R.E. WELSH. An adaptive nonlinear least-squares algorithm. *ACM Transactions on Math. Sci.*, 7:348-368, 1981.
- [5] J.E. DENNIS Jr., H.J. MARTINEZ and R.A. TAPIA. Convergence theory for the structured BFGS secant method with an application to nonlinear least squares. *JOTA*, 61(2):161-178, 1989.
- [6] J.E. DENNIS Jr. and R.B. SCHNABEL. *Numerical methods for unconstrained optimization and nonlinear equations*. Computational Mathematics. Prentice-Hall, New York, NY, 1983.
- [7] J.E. DENNIS Jr. and H. WOLKOWICZ. Sizing and least change secant methods. Report TR90-5, Department of Mathematical Sciences, Rice University, Houston, Tx 77251-1892, 1990.
- [8] J. HUSCHENS. On the use of product structure in secant methods for nonlinear least squares problems. Manuscript, Universität Trier, FB IV-Mathematik, Postfach 3825, 5500 Trier, Federal Republic of Germany, 1991.
- [9] R.M. McLEOD. Mean value theorems for vector valued functions. In *Proc. Ed. Math. Soc.*, volume 14, pages 197-209, 1964.
- [10] J.J. MORÉ, B.S. GARBOW, and K.E. HILLSTROM. Testing unconstrained optimization software. Technical Report ANL-78-324, Applied Mathematics Division, Argonne National Labs, Argonne, IL, 1978.
- [11] J. NOCEDAL and Y. YUAN. Analysis of self-scaling quasi-Newton method. Report, E.E.C.S. Department, Northwestern University, 1990, in preparation.
- [12] S.S. OREN and D.G. LUENBERGER. Self-scaling variable metric (ssvm) algorithms. *Man. Sci.*, 20:845-862, 1974.
- [13] S.S. OREN and E. SPEDICATO. Optimal conditioning of self-scaling variable metric algorithms. *Math. Progr.*, 10:70-90, 1976.
- [14] M. RAYDAN. On the Barzilai and Borwein choice of steplength for the gradient method. Technical Report TR90-11, Department of Mathematical Sciences, Rice University, Houston, Texas, 1990.
- [15] D.F. SHANNO and K. PHUA. Matrix conditioning and nonlinear optimization. *Math. Prog.*, 14:149-160, 1978.